

Interactive proofs with competing teams of no-signaling provers

Gus Gutoski

*Institute for Quantum Computing and School of Computer Science
University of Waterloo, Waterloo, Ontario, Canada*

December 3, 2010

(Minor revisions: October 14, 2011)

Abstract

This paper studies a generalization of multi-prover interactive proofs in which a verifier interacts with two competing teams of provers: one team attempts to convince the verifier to accept while the other attempts to convince the verifier to reject. Each team consists of two provers who jointly implement a *no-signaling* strategy. No-signaling strategies are a curious class of joint strategy that cannot in general be implemented without communication between the provers, yet cannot be used as a black box to establish communication between them. Attention is restricted in this paper to *two-turn* interactions in which the verifier asks questions of each of the four provers and decides whether to accept or reject based on their responses.

We prove that the complexity class of decision problems that admit two-turn interactive proofs with competing teams of no-signaling provers is a subset of PSPACE. This upper bound matches existing PSPACE lower bounds on the following two disparate and weaker classes of interactive proof:

1. Two-turn multi-prover interactive proofs with only one team of no-signaling provers.
2. Two-turn competing-prover interactive proofs with only one prover per team.

Our result implies that the complexity of these two models is unchanged by the addition of a second competing team of no-signaling provers in the first case and by the addition of a second no-signaling prover to each team in the second case. Moreover, our result unifies and subsumes prior PSPACE upper bounds on these classes.

1 Introduction

Interactive proofs were introduced in the mid-1980's as a generalization of the concept of efficient proof verification and the complexity class NP [Bab85, BM88, GMR89]. Informally speaking, an *interactive proof* is a conversation between a randomized polynomial-time *verifier* and a computationally unbounded *prover* regarding some common input string x . A decision problem L is said to admit an interactive proof if there exists a verifier such that (i) if x is a yes-instance of L then there is a prover who can convince the verifier to accept x with high probability, and (ii) if x is a no-instance of L then no prover can convince the verifier to accept x except with small probability. In a dramatic testament to the surprising power of randomization and interaction, it was soon discovered that every problem in PSPACE admits an interactive proof, yielding the well-known identity $IP = PSPACE$ [LFKN92, Sha92].

Multi-prover interactive proofs, no-signaling provers

The fruitful study of interactive proofs has prompted further generalization of the model. One such generalization is the *multi-prover* interactive proof model of Ben-Or *et al.* [BOGKW88] wherein several provers cooperate in their attempt to convince the verifier to accept the input string x . The key aspect that sets this model apart from single-prover interactive proofs is the fact that the provers cannot communicate with one another during the protocol. Amazingly, this small distinction is enough to increase the power of the model from PSPACE all the way up to NEXP [BFL91, FRS94], even when the interaction is restricted to only two turns with only two provers [FL92]. In terms of complexity classes, the corresponding identity is $\text{MIP} = \text{NEXP}$.

Intermediate classes of multi-prover interactive proofs are obtained by tinkering with the set of strategies available to the provers. Consider, for example, a joint strategy where the distribution of answers from one prover is independent of the question asked of the other prover—these are the *no-signaling* strategies. Clearly, such a strategy cannot be used in a black-box fashion by the provers to establish communication. At first glance it may seem that the no-signaling condition is equivalent to the standard definition of a multi-prover interactive proof. However, there exist no-signaling strategies that cannot be implemented without communication between the provers, suggesting that this model might be a nontrivial intermediary between single- and multi-prover interactive proofs.

Indeed, it was established by Ito, Kobayashi, and Matsumoto [IKM09] that the two-turn, two-prover protocol for PSPACE of Cai, Condon, and Lipton [CCL94] is sound even against no-signaling provers. By contrast, PSPACE is known not to admit two-turn *single-prover* interactive proofs unless the polynomial hierarchy collapses and $\text{PSPACE} = \text{AM}$ [Bab85, GS89]. A converse result was proven by Ito, who showed that every problem that admits a two-turn interactive proof with two no-signaling provers is also in PSPACE [Ito10]. Thus, the interactive proof model is even *more* sensitive to change than suggested by the difference between single- and multi-prover interactive proofs, as even the smaller difference between no-signaling and standard multi-prover interactive proofs is sufficient to make the jump from PSPACE up to NEXP (at least in the case of two turns and two provers).

In addition to this prior work, parallel repetition results for multi-prover interactive proofs with no-signaling provers were established in Refs. [Hol09, KR10]. The reader is referred to Ito [Ito10] for more detailed history and references.

Inspiration from quantum information

Though the present paper contains no formal discussion of quantum information, it is proper to acknowledge its role in motivating the study of no-signaling provers. Interest in this model was originally drawn from the study of multi-prover *quantum* interactive proofs, in which the provers (and possibly the verifier) are permitted to exchange and manipulate quantum information.

It is easy to see that interactive proofs with ordinary, “classical” provers are not affected by the ability of the provers to sample from a common source of randomness. Quantum provers, on the other hand, might use shared pieces of some entangled quantum state to implement a *nonlocal* strategy that correlates their messages in ways that cannot otherwise be achieved [Bel64]. (The phenomenon of nonlocality was famously branded by Einstein as “spooky action at a distance.”) Indeed, some classical protocols which are sound against classical provers are known to become unsound when the provers share entanglement [CHTW04, CGJ09].

Whereas the set of strategies that admit shared entanglement is highly complex, the set of no-signaling strategies is relatively simple and it includes entanglement-sharing strategies as a proper subset. So, for

example, any protocol that is sound against no-signaling provers is also sound against quantum provers who share entanglement. It is also interesting to find differences between no-signaling strategies and entanglement-sharing strategies, as this difference sheds light on the extent to which no-signaling can be used as a proxy for shared entanglement. In some protocols the allowance of arbitrary no-signaling strategies leads to implausible consequences [vD05, BBL⁺06]. Such protocols can be viewed as mathematical evidence against physical theories that admit so-called “super-strong” nonlocality such as that found in no-signaling strategies but not entanglement-sharing strategies. The present paper establishes a scenario in which two no-signalling provers are equivalent to two signaling provers.

Interactive proofs with competing provers

Another generalization of the single-prover model is an interactive proof with *competing provers*, in which one prover tries to convince the verifier to accept the input string x while the other prover tries to convince the verifier to reject x . One may consider proofs in which all messages are known to all provers (*complete* information) or in which each prover sees only the messages he exchanges with the verifier (*incomplete* information). These two forms of competing-prover interactive proofs were studied by several authors in the 1990’s [FST90, FS92, FKS95, FK97]. But for our purpose in this paper it only makes sense to consider protocols with incomplete information.

In the jargon of game theory, interactive proofs with competing provers are *zero-sum games*, about which there exists a vast body of literature in computer science, economics, and other disciplines. For instance, fast algorithms for zero-sum games of incomplete information in *extensive form* imply that the complexity class RG of problems that admit interactive proofs with competing provers is a subset of EXP [KM92, KMvS94]. Feige and Kilian proved the reverse containment [FK97], yielding the competing-prover analogy $RG = EXP$ of the aforementioned identity $IP = PSPACE$ for single-prover interactive proofs.

Feige and Kilian also studied *two-turn* interactive proofs with competing provers, providing a matching upper and lower bound of PSPACE on the complexity of this model [FK97]. The complexity of k -turn interactive proofs with competing provers for constants $k \geq 3$ is an open question of interest to both complexity theorists and game theorists alike.

Interactive proofs with competing teams of provers, our result

Multi-prover interactive proofs and interactive proofs with competing provers are two distinct generalizations of the single-prover model. The next logical step is to unify these two generalizations in the obvious way via interactive proofs with competing *teams* of provers. Combining established naming conventions for complexity classes based on interactive proofs, we let MRG denote the class of decision problems that admit interactive proofs with competing teams of provers.

To the author’s knowledge, this model was considered prior to the present work only by Feigenbaum, Koller, and Shor [FKS95]. Those authors studied this class under the game-theoretic guise of zero-sum games of *imperfect recall* and proved the containments

$$EXP^{NP} \subseteq MRG \subseteq \Sigma_2^{EXP} \cap \Pi_2^{EXP}$$

where Σ_2^{EXP} and Π_2^{EXP} are classes in the second level of the exponential hierarchy, which is the exponential-time version of the familiar polynomial hierarchy.

In this paper, we consider interactive proofs with competing teams of *no-signaling* provers. Our main result is as follows.

Theorem 1. *Every decision problem that admits a two-turn interactive proof with competing teams of two no-signaling provers per team is also in PSPACE.*

This upper bound matches the aforementioned PSPACE lower bounds on the following two disparate and weaker classes of interactive proof:

1. Two-turn multi-prover interactive proofs with only one team of no-signaling provers [CCL94, IKM09].
2. Two-turn competing-prover interactive proofs with only one prover per team [FK97].

Our result implies that the complexity of these two models is unchanged by the addition of a second competing team of no-signaling provers in the first case and by the addition of a second no-signaling prover to each team in the second case. Moreover, our result unifies and subsumes prior PSPACE upper bounds on these classes [Ito10, FK97].

Limitations of the present approach

Attention is restricted in this paper to interactions with no more than two no-signaling provers per team and no more than two messages exchanged with each prover. The purpose for this restriction, quite simply, is that this class of interactions appears to be the largest to which our techniques apply.

For all we know, interactions with three messages for a prover or three provers on a team could be sufficiently powerful to capture all of EXP. Indeed, it is consistent with current knowledge that a three-message protocol for EXP might require only *one* prover per team, or that a three-prover no-signaling protocol for EXP might require only *one* team of provers. Given this paucity of upper bounds for similar, seemingly weaker models it is hoped that any reservation at the restrictions in our model is more than compensated by the fact that we are able to say anything at all about it.

Let us list some natural extensions of the two-prover, two-turn model and point out exactly where our method fails for these extensions.

More than two turns, only one prover per team. Perhaps the most important open problem related to our work is the complexity of k -turn interactive proofs with competing provers for constants $k \geq 3$. This problem, which dates back at least to 1997 [FK97], is still open even in the special case of only one prover per team. With only one prover per team, the question is really a game-theoretic question with a much wider application than just interactive proofs.

Our method fails for this case because we do not have a bound on the verifier matrix of the form $V \leq e_{\mathcal{A}_{01}\mathcal{B}_{01}} p^*$ such as that appearing in Proposition 3. Thus, we do not obtain a good enough bound on the loss vectors appearing in our variant of the multiplicative weights update method.

More than two turns, only one team of no-signaling provers. The complexity of k -turn multi-prover interactive proofs with two no-signaling provers is still open for $k \geq 3$, even with only one team of provers [Ito10]. For ordinary multi-prover interactive proofs—in which the provers are not allowed to implement arbitrary no-signaling strategies—it is known that a multi-turn protocol with any number of provers can be simulated by another protocol with only two turns and two provers [FL92].

Our method fails here for the same reason as above—that we cannot bound the loss vectors in the multiplicative weights update method for a multi-turn verifier.

More than two provers, only one team of no-signaling provers. Similarly, the complexity of two-turn multi-prover interactive proofs with more than two no-signaling provers is still open, even with only one

team of provers [Ito10]. As mentioned above, ordinary multi-prover interactive proofs require only two provers [FL92].

Our method does not extend to this case either, as there is no known analogue of Lemma 6 for more than two provers.

Quantum verifier and/or provers. Even with two no-signaling provers, two turns of interaction, and only one team of provers, it is still not known that the PSPACE upper bound holds when either the verifier or provers can send quantum messages [Ito10]. Here the problem is that Lemma 6 does not hold for quantum states.

Techniques

Theorem 1 is proven by means of an efficient parallel algorithm that, given an explicit description of a verifier and an accuracy parameter δ , finds no-signaling strategies for the teams that are within δ of optimal. Containment in PSPACE then follows in the usual way by observing that the description of the verifier has size exponential in the length of the input string x and then employing the fact that a parallel algorithm with succinct input can be simulated in polynomial space [Bor77].

Our algorithm is an example of the *multiplicative weights update method (MWUM)* as discussed in the survey paper [AHK05] and in the PhD thesis of Kale [Kal07]. (See also Ref. [WK06].) In its simplest form, the MWUM solves a min-max optimization problem on probability distributions. In the present paper we use the MWUM to optimize not just a *single* distribution, but many distributions *simultaneously* in the form of a stochastic matrix that represents a strategy for one of the teams. This trick seems to work only for two-turn protocols, as otherwise it is not clear how to ensure sufficient accuracy.

Let us compare our algorithm to the two previous algorithms it subsumes:

- The polynomial-space algorithm of Feige and Kilian for two-turn interactive proofs with competing provers [FK97] is a complicated and highly specialized precursor to the MWUM that, like our algorithm, optimizes over stochastic matrices that represent strategies for the provers.

Their algorithm works by nondeterministically guessing the entries of the matrix and scanning them in a read-once fashion. This approach cannot be extended to optimize over no-signaling strategies, as the read-once model does not allow verification of the no-signaling condition.

- The parallel algorithm of Ito for two-turn, two-prover interactive proofs with no-signaling provers [Ito10] is essentially a reduction to the *mixed packing and covering problem*, which is a special type of linear program that is known to admit an efficient parallel algorithm [You01].

This approach, too, cannot be extended to competing teams of no-signaling provers, as any linear programming formulation of the protocol is unlikely to be a mixed packing and covering problem.

Our study has benefitted from the valuable experience of recent applications of the MWUM to parallel algorithms for quantum complexity classes [JW09, JUW09, JJUW10, Wu10, GW11]. Indeed, we follow the same high-level approach as the recent proof of $\text{DQIP} = \text{DIP} = \text{PSPACE}$ [GW11]. Namely,

- The domain of admissible (no-signaling) strategies is a strict subset of the “natural” domain (stochastic matrices) for the MWUM.
- To get around this problem, the strategy domain is extended to *all* the stochastic matrices and a *penalty term* is introduced so as to remove any incentive for a team to use an inadmissible strategy. (See Section 3).




Figure 1: A two-turn interactive proof with competing teams of two no-signaling provers per team.

- Finally, one must prove a “rounding” theorem (Corollary 4.1), which establishes that near-optimal, fully-admissible strategies can be obtained from near-optimal strategies in the extended domain with penalty term.

2 Preliminaries

2.1 Definition of two-turn interactive proofs with competing teams of provers

In this paper we are concerned with decision problems that admit two-turn interactive proofs with competing teams of no-signaling provers. Let us clarify this concept. A *two-turn verifier* is a randomized polynomial-time algorithm that, given an input string x , produces questions i, j for the two teams of provers. The teams select their answers k, l (possibly using randomness to do so) and then the verifier accepts or rejects the input x according to some boolean function of i, j, k, l . For convenience, the teams shall be called *Team Alice* and *Team Bob*. It is the goal of Team Alice to convince the verifier to accept the input string x , while Team Bob’s goal is to convince the verifier to reject x .

In the protocols we consider each team consists of two provers. The provers of Team Alice shall be called $Alice_0$ and $Alice_1$, while the provers of Team Bob shall be called Bob_0 and Bob_1 . Each individual prover on each team receives his or her own private question and supplies his or her own separate answer to the verifier. In particular, the question i asked of Team Alice is actually a pair $i = (i_0, i_1)$ with question i_c going to prover $Alice_c$ for both values of the bit $c \in \{0, 1\}$. Similarly, the question j asked of Team Bob is also a pair $j = (j_0, j_1)$ with question j_c going to prover Bob_c . The answers k, l received from the two teams are also pairs $k = (k_0, k_1)$ and $l = (l_0, l_1)$ with answers k_c and l_c coming from $Alice_c$ and Bob_c , respectively. The entire interaction is illustrated in Figure 1.

Each team may jointly implement any no-signaling strategy in order to produce its answers. Briefly, a strategy for, say, Team Alice is *no-signaling* if the marginal distribution on answers k_0 from $Alice_0$ does not depend upon the question i_1 asked of $Alice_1$ and *vice versa*. No-signaling strategies are discussed in greater detail in Section 2.5.

A decision problem L is said to admit a two-turn interactive proof with competing teams of no-signaling provers with *completeness* c and *soundness* s if there exists a fixed two-turn verifier with the following properties:

Completeness. If the input string x is a yes-instance of L then there exists a no-signaling strategy for Team

Alice that convinces the verifier to accept x with probability at least c , regardless of the no-signaling strategy employed by Team Bob.

Soundness. If the input string x is a no-instance of L then there exists a no-signaling strategy for Team Bob that convinces the verifier to reject x with probability at least $1 - s$, regardless of the no-signaling strategy employed by Team Alice.

The completeness and soundness parameters need not be fixed constants. Rather, they may vary as a function of the input string x . The complexity class $\text{MRG}_{\text{ns}}(2, 2)$ consists of all decision problems that admit two-turn interactive proofs with competing teams of two no-signaling provers per team with completeness c and soundness s such that there exists a fixed polynomial-bounded function p on strings with $c - s \geq 1/p$. (The first parameter of the class $\text{MRG}_{\text{ns}}(2, 2)$ denotes the number of provers per team, the second denotes the number of turns in the protocol. It is also common to parameterize interactive proof classes according to the number of *rounds* of communication, rather than the number of *turns*. Under this scheme, the class $\text{MRG}_{\text{ns}}(2, 2)$ might be called $\text{MRG}_{\text{ns}}(2, 1)$ by some authors.)

In this paper we prove $\text{MRG}_{\text{ns}}(2, 2) \subseteq \text{PSPACE}$. It then follows from existing lower bounds on weaker classes [IKM09, FK97] that

$$\text{MRG}_{\text{ns}}(2, 2) = \text{PSPACE}.$$

2.2 Notation, the Kronecker product

To each interactive proof with input x we associate eight distinct finite-dimensional real Euclidean spaces—four *question* spaces and four *answer* spaces. These spaces are denoted as follows for both $c \in \{0, 1\}$:

\mathcal{S}_c	The question space for prover Alice _{c}	\mathcal{A}_c	The answer space for prover Alice _{c}
\mathcal{T}_c	The question space for prover Bob _{c}	\mathcal{B}_c	The answer space for prover Bob _{c}

The dimension of each space is the number of distinct questions or answers available to that prover. (For example, prover Alice₀ can be asked any of $\dim(\mathcal{S}_0)$ distinct questions and may respond with any of $\dim(\mathcal{A}_0)$ distinct answers.) Individual questions or answers are indexed by positive integers denoted for both $c \in \{0, 1\}$ as follows:

$$\begin{aligned} \text{Questions for Alice}_c : & i_c = 1, \dots, \dim(\mathcal{S}_c) \\ \text{Questions for Bob}_c : & j_c = 1, \dots, \dim(\mathcal{T}_c) \\ \text{Answers from Alice}_c : & k_c = 1, \dots, \dim(\mathcal{A}_c) \\ \text{Answers from Bob}_c : & l_c = 1, \dots, \dim(\mathcal{B}_c) \end{aligned}$$

Since the verifier acts in polynomial time, the bit length of the questions and answers is at most a polynomial in the bit length $|x|$ of the input string x . Since n bits suffice to encode 2^n distinct questions or answers, the dimension of the spaces $\mathcal{S}_c, \mathcal{T}_c, \mathcal{A}_c, \mathcal{B}_c$ can be exponential in $|x|$.

The *Kronecker product* (or *tensor product*) of two spaces \mathcal{X}, \mathcal{Y} is another space with dimension $\dim(\mathcal{X}) \dim(\mathcal{Y})$. This product space is typically denoted by $\mathcal{X} \otimes \mathcal{Y}$, which we abbreviate to $\mathcal{X}\mathcal{Y}$. Kronecker products involving the eight spaces $\mathcal{S}_c, \mathcal{T}_c, \mathcal{A}_c, \mathcal{B}_c$ are further abbreviated so that

$$\mathcal{S}_{01} = \mathcal{S}_0 \mathcal{S}_1 = \mathcal{S}_0 \otimes \mathcal{S}_1$$

and so on. The Kronecker product extends in a natural way to vectors and linear operators. In this paper each vector or linear operator is implicitly associated with its representation as a column or a matrix, for

which the Kronecker product is given by a straightforward formula. For example, if A, B are 2×2 matrices given by

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}, \quad B = \begin{bmatrix} p & q \\ r & s \end{bmatrix}$$

then the Kronecker product $A \otimes B$ is given by

$$A \otimes B = \begin{bmatrix} aB & bB \\ cB & dB \end{bmatrix} = \begin{bmatrix} a \begin{bmatrix} p & q \\ r & s \end{bmatrix} & b \begin{bmatrix} p & q \\ r & s \end{bmatrix} \\ c \begin{bmatrix} p & q \\ r & s \end{bmatrix} & d \begin{bmatrix} p & q \\ r & s \end{bmatrix} \end{bmatrix} = \begin{bmatrix} ap & aq & bp & bq \\ ar & as & br & bs \\ cp & cq & dp & dq \\ cr & cs & dr & ds \end{bmatrix}.$$

This definition extends in the obvious way to arbitrary matrices of any dimension, including column vectors and other non-square matrices.

We also make use of the following symbols:

- $e_{\mathcal{X}}$ The all-ones vector of dimension $\dim(\mathcal{X})$.
- $I_{\mathcal{X}}$ The identity matrix acting on \mathcal{X} .
- M^* The *adjoint* of a linear mapping M . If M is a matrix or column vector then M^* is simply the transpose of M .
- $\langle A, B \rangle$ The *matrix inner product*, defined as $\text{Tr}(A^*B)$. This inner product is defined only when the dimensions of A, B are equal. If A, B are vectors then $\langle A, B \rangle$ is called the *vector inner product*.
- \leq, \geq Matrix inequalities are entrywise.
- \bar{c} Given a bit $c \in \{0, 1\}$, the compliment \bar{c} is given by $\bar{c} = 1$ if $c = 0$, otherwise $\bar{c} = 0$.

2.3 Min-max formalism for interactive proofs with competing provers

Given a fixed two-turn verifier and a fixed input string x , let $\pi_{i,j}$ denote the probability with which the verifier asks questions $i = (i_0, i_1)$ to Team Alice and $j = (j_0, j_1)$ to Team Bob. For each 4-tuple (i, j) of questions to the provers let $v_{i,j} \in \mathcal{A}_{01}\mathcal{B}_{01}$ denote the 0-1 vector of *payouts* to Team Bob. That is, for each $k = (k_0, k_1)$ and each $l = (l_0, l_1)$ the (k, l) th entry of $v_{i,j}$ is either zero or one according to whether the verifier accepts or rejects x in the event that the verifier asks questions (i, j) to the teams and they respond with answers (k, l) .^{1,2} Consider the entrywise nonnegative matrix

$$V : \mathcal{S}_{01}\mathcal{T}_{01} \rightarrow \mathcal{A}_{01}\mathcal{B}_{01}$$

whose (i, j) th column is $\pi_{i,j}v_{i,j}$. This matrix uniquely specifies the actions of the verifier.

Strategies for the teams are specified as follows. For each pair i of questions let $a_i \in \mathcal{A}_{01}$ denote the probability vector of Team Alice's responses to i . That is, for each pair k of answers the k th entry of a_i denotes the probability with which Team Alice replies with answers k given that questions i were asked. Thus, the actions of Team Alice are uniquely specified by the stochastic matrix

$$A : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$$

¹One could consider a more general referee in which the payouts are awarded probabilistically so that each entry of $v_{i,j}$ lies in the interval $[0, 1]$. But it is easily seen that this model is equivalent to the one we have just described.

²The payout vector $v_{i,j}$ is defined so that 0 indicates acceptance of x while 1 indicates rejection. This arbitrary choice is opposite of convention, but it better facilitates the forthcoming presentation of our multiplicative weights update algorithm.

whose i th column is a_i . Similarly, for each pair j of questions let $b_j \in \mathcal{B}_{01}$ denote the probability vector of Team Bob's responses to j . The actions of Team Bob are uniquely specified by the stochastic matrix

$$B : \mathcal{T}_{01} \rightarrow \mathcal{B}_{01}$$

whose j th column is b_j . Not every stochastic matrix denotes a valid no-signaling strategy for the teams. Criteria for no-signaling strategies are discussed in Section 2.5. For now, it suffices to note that the set of all strategies available to each team is a compact convex subset of stochastic matrices.

Conditioned on the verifier asking questions (i, j) , it is clear that the probability of rejection is given by the vector inner product

$$\langle v_{i,j}, a_i \otimes b_j \rangle.$$

It follows that the probability of rejection—taken over all questions (i, j) —given strategies A for Team Alice and B for Team Bob is given by the matrix inner product

$$\Pr[V \text{ rejects } x \mid A, B] = \langle V, A \otimes B \rangle = \sum_{i,j} \pi_{i,j} \langle v_{i,j}, a_i \otimes b_j \rangle.$$

Of course, Team Bob wishes to maximize this quantity while Team Alice wishes to minimize this quantity. Given that the above inner product is bilinear in (A, B) and that the sets of admissible strategies for the two teams are compact and convex, it follows from standard min-max theorems [Vil38, Fan53] that every interactive proof with verifier V has an *equilibrium value*, which we denote by $\lambda(V)$, given by

$$\lambda(V) = \min_A \max_B \langle V, A \otimes B \rangle = \max_B \min_A \langle V, A \otimes B \rangle$$

where the minimum is over all no-signaling matrices $A : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$ and the maximum is over all no-signaling matrices $B : \mathcal{T}_{01} \rightarrow \mathcal{B}_{01}$. In particular, for every protocol there exists at least one *equilibrium point* (A^*, B^*) with the property that

$$\begin{aligned} \langle V, A^* \otimes B \rangle &\leq \lambda(V) \quad \text{for all } B, \\ \langle V, A \otimes B^* \rangle &\geq \lambda(V) \quad \text{for all } A. \end{aligned}$$

Thus, the strategy B^* always ensures maximum likelihood of rejection, while A^* always ensures minimum likelihood of rejection.

This min-max theorem applies to every min-max expression considered throughout this paper. Henceforth we do not bother to explicitly remark upon this fact. Here and throughout the paper we adopt the convention that for any min-max problem of the form

$$\nu(g) = \min_{a \in \mathbf{A}} \max_{b \in \mathbf{B}} g(a, b)$$

elements $\tilde{a} \in \mathbf{A}$ and $\tilde{b} \in \mathbf{B}$ are δ -*optimal* if

$$\begin{aligned} g(\tilde{a}, b) &\leq \nu(g) + \delta \quad \text{for all } b \in \mathbf{B}, \\ g(a, \tilde{b}) &\geq \nu(g) - \delta \quad \text{for all } a \in \mathbf{A}. \end{aligned}$$

Elements that are 0-optimal—such as A^*, B^* above—are simply called *optimal*.

2.4 Notation for marginal distributions

Before we discuss no-signaling strategies in detail it is beneficial to introduce notation for marginal probability distributions that will be used throughout the remainder of this paper. Suppose, for instance, that $a \in \mathcal{A}_{01}$ is a probability vector of answers from Team Alice to some question from the verifier. We let $\text{mar}_{\mathcal{A}_1}(a) \in \mathcal{A}_0$ denote the probability vector for the marginal distribution on answers from the prover Alice₀. Basic probability theory dictates that the mapping $\text{mar}_{\mathcal{A}_1}$ satisfy

$$k_0\text{th entry of } \text{mar}_{\mathcal{A}_1}(a) \equiv \sum_{k_1=1}^{\dim(\mathcal{A}_1)} (k_0, k_1)\text{th entry of } a.$$

Of course, this mapping may be extended to arbitrary real vectors. For arbitrary spaces \mathcal{X}, \mathcal{Y} the linear mapping $\text{mar}_{\mathcal{Y}}$ is defined by

$$\text{mar}_{\mathcal{Y}} : \mathcal{X}\mathcal{Y} \rightarrow \mathcal{X} : x \otimes y \mapsto \langle e_{\mathcal{Y}}, y \rangle x.$$

(The matrix representation of $\text{mar}_{\mathcal{Y}}$ is $e_{\mathcal{Y}}^* \otimes I_{\mathcal{X}}$.) While this mapping is primarily intended to denote marginal probability distributions, we will have occasion to use it on non-probability vectors in this paper.

The mapping $\text{mar}_{\mathcal{Y}}$ is to vectors as the *partial trace* is to square matrices. Readers familiar with quantum information know that the state of a quantum register can be computed from a joint state of several registers via the partial trace. So too with probability distributions: the distribution on states of a classical register can be computed from a joint distribution on states of several registers via $\text{mar}_{\mathcal{Y}}$.

The mapping $\text{mar}_{\mathcal{Y}}$ extends naturally from vectors to matrices by applying $\text{mar}_{\mathcal{Y}}$ to each column:

$$i\text{th column of } \text{mar}_{\mathcal{Y}}(A) \equiv \text{mar}_{\mathcal{Y}}(i\text{th column of } A).$$

So, for example, if Team Alice acts according to the stochastic matrix A then the stochastic matrix

$$\text{mar}_{\mathcal{A}_1}(A) : \mathcal{S}_{01} \rightarrow \mathcal{A}_0$$

describes the “marginal” strategy for prover Alice₀. That is, the (i_0, i_1) th column of $\text{mar}_{\mathcal{A}_1}(A)$ is the distribution on answers k_0 from Alice₀ given questions (i_0, i_1) from the verifier.

2.5 Characterization of no-signaling strategies

Recall that a strategy for Team Alice is *no-signaling* if for both values of the bit $c \in \{0, 1\}$ the marginal distribution on answers k_c from Alice_c does not depend on the question $i_{\bar{c}}$ asked of Alice _{\bar{c}} .

In terms of Team Alice’s stochastic matrix A , this condition means that for each i_c the (i_0, i_1) th column of $\text{mar}_{\mathcal{A}_{\bar{c}}}(A)$ is identical for all subindices $i_{\bar{c}}$. Letting a_{i_c} denote this fixed probability vector and letting $A_c : \mathcal{S}_c \rightarrow \mathcal{A}_c$ denote the stochastic matrix whose columns are a_{i_c} , the above condition can be written as

$$\text{mar}_{\mathcal{A}_c}(A) = A_c \otimes e_{\mathcal{S}_{\bar{c}}}^*.$$

We have just proven the following simple proposition.

Proposition 2 (Characterization of no-signaling strategies). *A stochastic matrix $A : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$ denotes a no-signaling strategy for Team Alice if and only if for both values of the bit $c \in \{0, 1\}$ there exists a stochastic matrix $A_c : \mathcal{S}_c \rightarrow \mathcal{A}_c$ such that*

$$\text{mar}_{\mathcal{A}_{\bar{c}}}(A) = A_c \otimes e_{\mathcal{S}_{\bar{c}}}^*.$$

A similar characterization holds for Team Bob.

Stochastic matrices A meeting this condition are called *no-signaling matrices*. The matrices A_c are said to *witness* the fact that A is a no-signaling matrix. It follows immediately from Proposition 2 that the set of all no-signaling strategies available to each team is compact and convex—a fact already used in Section 2.3 to assert the existence of optimal strategies for the teams.

3 A relaxed min-max problem with penalties

As mentioned in the introduction, the MWUM in its simplest form solves min-max optimization problems over probability vectors. We optimize over stochastic matrices for the teams by using the MWUM simultaneously on each column of these matrices—a trick that works only for two-turn protocols, as we shall soon see.

We noted in Section 2.5 that the no-signaling matrices available to the teams form a strict subset of the stochastic matrices. In order to optimize only over no-signaling matrices, in this section we specify a new min-max optimization problem $\mu(V)$ in which the teams may use *arbitrary* strategies but pay a *penalty* for strategies that violate the no-signaling condition. By a careful choice of penalty, we remove the incentive of the teams to select inadmissible strategies without ruining the precarious convergence properties of the MWUM.

Some preliminary observations are given in Section 3.1 before the formal definition of the new min-max problem $\mu(V)$ in Section 3.2. Equivalence of $\mu(V)$ and $\lambda(V)$ is proven in Section 3.3 with proofs of some lemmas in Section 3.4.

3.1 Bounds on two-turn verifiers

First, for ease of notation we let Φ_V denote the unique linear transformation satisfying

$$\langle V, A \otimes B \rangle = \langle \Phi_V(A), B \rangle = \langle A, \Phi_V^*(B) \rangle$$

for all matrices A, B . Though a precise formula for Φ_V is of little use in this paper, for completeness we note that

$$\begin{aligned}\Phi_V(A) &= \text{Tr}_{\mathcal{S}_{01}} ((A^* \otimes I_{\mathcal{B}_{01}}) V) \\ \Phi_V^*(B) &= \text{Tr}_{\mathcal{T}_{01}} ((I_{\mathcal{A}_{01}} \otimes B^*) V)\end{aligned}$$

where $\text{Tr}_{\mathcal{S}_{01}}$ and $\text{Tr}_{\mathcal{T}_{01}}$ denote *partial trace* transformations. At the risk of hijacking terminology from functional analysis, the matrix $\Phi_V(A)$ can be viewed as a *partial inner product* between V and A . This matrix can also be viewed as a new two-turn verifier for Team Bob obtained by “hard-wiring” Team Alice’s strategy A into the original verifier V .

Next, let $p \in \mathcal{S}_{01} \mathcal{T}_{01}$ denote the probability vector for the distribution on questions asked by the verifier. In the notation of Section 2.3, the (i, j) th entry of p is $\pi_{i,j}$ —the probability with which the verifier asks questions i to Team Alice and j to Team Bob. Let $p_{\text{Alice}} \in \mathcal{S}_{01}$ denote the marginal distribution

$$p_{\text{Alice}} = \text{mar}_{\mathcal{T}_{01}}(p)$$

on questions to Team Alice, so that the i th entry of p_{Alice} is $\sum_j \pi_{i,j}$. It is not hard to see that

$$V \leq e_{\mathcal{A}_{01} \mathcal{B}_{01}} p^*$$

with equality achieved in the extreme case that each of the verifier's payout vectors $v_{i,j}$ is equal to the all-ones vector $e_{\mathcal{A}_{01}\mathcal{B}_{01}}$. (Recall that matrix inequalities are entrywise.) Similarly, it is easy to prove analogous inequalities for $\Phi_V(A)$, $\Phi_V^*(B)$. For example:

Proposition 3. *For any stochastic matrix $B : \mathcal{T}_{01} \rightarrow \mathcal{B}_{01}$ it holds that $\Phi_V^*(B) \leq e_{\mathcal{A}_{01}} p_{\text{Alice}}^*$.*

Proof. Let $A : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$ be any nonnegative matrix and let a_i, b_j denote the columns of A, B , respectively. Then

$$\langle A, \Phi_V^*(B) \rangle = \langle V, A \otimes B \rangle \leq \langle e_{\mathcal{A}_{01}\mathcal{B}_{01}} p^*, A \otimes B \rangle = \sum_{i,j} \pi_{i,j} \langle e_{\mathcal{A}_{01}}, a_i \rangle \langle e_{\mathcal{B}_{01}}, b_j \rangle$$

As B is stochastic it must be that $\langle e_{\mathcal{B}_{01}}, b_j \rangle = 1$ for each j . The above expression then simplifies to

$$\sum_i \left(\sum_j \pi_{i,j} \right) \langle e_{\mathcal{A}_{01}}, a_i \rangle = \langle e_{\mathcal{A}_{01}} p_{\text{Alice}}^*, A \rangle.$$

As this inequality holds for all nonnegative matrices A it must be that $\Phi_V^*(B) \leq e_{\mathcal{A}_{01}} p_{\text{Alice}}^*$ as claimed. \square

3.2 Definition of the relaxed min-max problem

The relaxation $\mu(V)$ of $\lambda(V)$ is defined by

$$\mu(V) = \min_{(A, A_0, A_1)} \max_{(B, \Pi_0, \Pi_1)} \langle f_V(A, A_0, A_1), (B, \Pi_0, \Pi_1) \rangle$$

where the triples (A, A_0, A_1) and (B, Π_0, Π_1) have the form

$$\begin{array}{lll} A : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01} & \text{any stochastic} & \\ A_c : \mathcal{S}_c \rightarrow \mathcal{A}_c & \text{any stochastic} & c \in \{0, 1\} \\ B : \mathcal{T}_{01} \rightarrow \mathcal{B}_{01} & \text{no-signaling only} & \\ \Pi_c : \mathcal{S}_{01} \rightarrow \mathcal{A}_c & 0 \leq \Pi_c \leq e_{\mathcal{A}_c} p_{\text{Alice}}^* & c \in \{0, 1\}. \end{array}$$

The linear mapping f_V appearing in the inner product (and its adjoint) is defined by

$$\begin{aligned} f_V : (A, A_0, A_1) &\mapsto (\Phi_V(A), \text{mar}_{\mathcal{A}_1}(A) - A_0 \otimes e_{\mathcal{S}_1}^*, \text{mar}_{\mathcal{A}_0}(A) - A_1 \otimes e_{\mathcal{S}_0}^*) \\ f_V^* : (B, \Pi_0, \Pi_1) &\mapsto (\Phi_V^*(B) + e_{\mathcal{A}_1} \otimes \Pi_0 + e_{\mathcal{A}_0} \otimes \Pi_1, -\Pi_0(I_{\mathcal{S}_0} \otimes e_{\mathcal{S}_1}), -\Pi_1(I_{\mathcal{S}_1} \otimes e_{\mathcal{S}_0})) \end{aligned}$$

so that

$$\langle f_V(A, A_0, A_1), (B, \Pi_0, \Pi_1) \rangle = \langle V, A \otimes B \rangle + \sum_{c \in \{0, 1\}} \langle \text{mar}_{\mathcal{A}_c}(A) - A_c \otimes e_{\mathcal{S}_c}^*, \Pi_c \rangle$$

for all (A, A_0, A_1) and all (B, Π_0, Π_1) . (The adjoint mapping f_V^* is not used until the algorithm of Figure 2 and its proof of correctness in Proposition 8.)

Intuition

Some explanation is in order. As with the original min-max problem $\lambda(V)$, the matrices A and B represent the strategies employed by the teams. Note, however, that in the definition of $\mu(V)$ Team Alice is now free to choose among arbitrary stochastic matrices for its strategy. The matrices A_0, A_1 for Team Alice are purported witnesses to the claim that A is a valid no-signaling matrix.

For the moment, we are concerned with relaxing the domain only of Team Alice's strategies, so Bob's strategy B must still be no-signaling. Bob's strategies will be addressed in Section 4.2. The matrices Π_0, Π_1 for Team Bob are *penalty matrices*—they are the means by which Team Bob penalizes Team Alice according to the extent that A_0, A_1 are false witnesses to the claim that A is no-signaling.

The new objective function $\langle f_V(A, A_0, A_1), (B, \Pi_0, \Pi_1) \rangle$ equals the old objective function $\langle V, A \otimes B \rangle$ plus two *penalty terms*. If A is not a no-signaling matrix then the difference matrix

$$\Delta_c \equiv \text{mar}_{\mathcal{A}_c}(A) - A_c \otimes e_{\mathcal{S}_c}^*$$

must be nonzero for at least one c . In this case, Bob selects Π_c to pick out the positive entries of Δ_c , which are then added the verifier's probability of rejection.

Let us informally explain why the restriction $0 \leq \Pi_c \leq e_{\mathcal{A}_c} p_{\text{Alice}}^*$ on penalty matrices is sufficient to remove Team Alice's incentive to cheat. Suppose the k_c th entry of the i th column of the difference matrix Δ_c is a positive real number $\delta > 0$ and suppose that A' is a valid no-signaling matrix witnessed by A_0, A_1 . Since the verifier asks questions i of Team Alice with probability π_i , it must be that, when selecting the probability with which to answer k_c , the advantage gained by Team Alice from using the inadmissible strategy A instead of the no-signaling strategy A' is at most $\delta\pi_i$. By selecting a penalty matrix Π_c so that the k_c th entry of the i th column of Π_c is equal to π_i , Team Bob adds precisely the quantity $\delta\pi_i$ to the verifier's probability of rejection, thus eliminating the advantage obtained by Team Alice in acting according to A instead of A' for this particular choice of questions i and answer k_c from Alice _{c} .

Repeating this logic for all entries (i, k_c) of Δ_c , we find that Team Bob should select the penalty matrix Π_c so that the (i, k_c) th entry is either zero or π_i according to whether the corresponding entry of Δ_c is nonpositive or positive. A penalty matrix of this form is called *optimal for* (A, A_0, A_1) and satisfies

$$\langle \Delta_c, \Pi_c \rangle = \langle \Delta_c^+, e_{\mathcal{A}_c} p_{\text{Alice}}^* \rangle$$

where Δ_c^+ is the positive part of Δ_c . (Here the *positive part* of a real matrix X is the matrix X^+ with the property that if x is any entry of X then the corresponding entry of X^+ is $\max\{0, x\}$.)

3.3 Equivalence of the two min-max problems

We are now ready to prove the desired “rounding theorem” mentioned in the introduction, a corollary of which is the equivalence of the min-max problems $\mu(V)$ and $\lambda(V)$ (Corollary 4.1). The theorem employs two lemmas and their corollaries, the proofs of which appear below in Section 3.4.

Theorem 4 (Rounding theorem). *Let (A, A_0, A_1) be a feasible solution for $\mu(V)$ and let Π_0^A, Π_1^A be optimal penalties for (A, A_0, A_1) . There exists a no-signaling matrix A_{ns} witnessed by A_0, A_1 such that for all stochastic matrices B it holds that*

$$\langle V, A_{\text{ns}} \otimes B \rangle \leq \langle f_V(A, A_0, A_1), (B, \Pi_0^A, \Pi_1^A) \rangle.$$

Moreover, A_{ns} can be computed efficiently in parallel given (A, A_0, A_1) .

Proof. For both $c \in \{0, 1\}$ let Δ_c^+ be the positive part of $\text{mar}_{\mathcal{A}_c}(A) - A_c \otimes e_{\mathcal{S}_c}^*$ and observe that

$$\Delta_c^+ \leq \text{mar}_{\mathcal{A}_c}(A).$$

By Corollary 5.1 below there exists a preimage $D_0^+ \geq 0$ of Δ_0^+ with

$$\begin{aligned} A - D_0^+ &\geq 0 \\ \text{mar}_{\mathcal{A}_1}(D_0^+) &= \Delta_0^+. \end{aligned}$$

Let Γ_1^+ be the positive part of $\text{mar}_{\mathcal{A}_0}(A - D_0^+) - A_1 \otimes e_{\mathcal{S}_0}^*$. As with Δ_c above, observe that

$$\Gamma_1^+ \leq \text{mar}_{\mathcal{A}_0}(A - D_0^+).$$

(Moreover, it is easy to see that $\Gamma_1^+ \leq \Delta_1^+$ —a fact we employ later in this proof.) Apply Corollary 5.1 again to obtain a preimage $C_1^+ \geq 0$ of Γ_1^+ with

$$\begin{aligned} A - D_0^+ - C_1^+ &\geq 0 \\ \text{mar}_{\mathcal{A}_0}(C_1^+) &= \Gamma_1^+. \end{aligned}$$

Thus, we have a matrix $A - D_0^+ - C_1^+ \geq 0$ such that for both $c \in \{0, 1\}$ it holds that

$$\text{mar}_{\mathcal{A}_c}(A - D_0^+ - C_1^+) \leq A_c \otimes e_{\mathcal{S}_c}^*.$$

Hence there exist nonnegative matrices $T_c : \mathcal{S}_{01} \rightarrow \mathcal{A}_c$ with

$$\text{mar}_{\mathcal{A}_c}(A - D_0^+ - C_1^+) + T_c = A_c \otimes e_{\mathcal{S}_c}^*.$$

Applying $\text{mar}_{\mathcal{A}_c}$ to both sides of this equation we see that $\text{mar}_{\mathcal{A}_0}(T_0) = \text{mar}_{\mathcal{A}_1}(T_1)$. By Corollary 6.1 below there exists a nonnegative matrix $T : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$ with $\text{mar}_{\mathcal{A}_c}(T) = T_c$ for both $c \in \{0, 1\}$. The desired no-signaling matrix A_{ns} is given by

$$A_{\text{ns}} = A - D_0^+ - C_1^+ + T.$$

As D_0^+ , C_1^+ , and T can be computed efficiently in parallel, so too can A_{ns} . To see that A_{ns} is a no-signaling matrix witnessed by A_0, A_1 it suffices to observe that

$$\text{mar}_{\mathcal{A}_c}(A_{\text{ns}}) = \text{mar}_{\mathcal{A}_c}(A - D_0^+ - C_1^+) + T_c = A_c \otimes e_{\mathcal{S}_c}^*.$$

It remains only to verify the stated inequality. To this end, we have

$$\begin{aligned} \langle V, A_{\text{ns}} \otimes B \rangle &= \langle A, \Phi_V^*(B) \rangle - \langle D_0^+ + C_1^+, \Phi_V^*(B) \rangle + \langle T, \Phi_V^*(B) \rangle \\ &\leq \langle A, \Phi_V^*(B) \rangle + \langle T, \Phi_V^*(B) \rangle \\ &\leq \langle A, \Phi_V^*(B) \rangle + \langle T, e_{\mathcal{A}_{01}} p_{\text{Alice}}^* \rangle \end{aligned}$$

As A_{ns} and A are both stochastic matrices, it must be that $D_0^+ + C_1^+$ and T have the same column sums. As $\langle T, e_{\mathcal{A}_{01}} p_{\text{Alice}}^* \rangle$ equals the sum of the column sums of T weighted according to p_{Alice} , the matrix T can be replaced by $D_0^+ + C_1^+$ without affecting this inner product. That is

$$\langle T, e_{\mathcal{A}_{01}} p_{\text{Alice}}^* \rangle = \langle D_0^+ + C_1^+, e_{\mathcal{A}_{01}} p_{\text{Alice}}^* \rangle.$$

Expanding the right side of this equality we obtain

$$\langle \text{mar}_{\mathcal{A}_1}(D_0^+), e_{\mathcal{A}_0} p_{\text{Alice}}^* \rangle + \langle \text{mar}_{\mathcal{A}_0}(C_1^+), e_{\mathcal{A}_1} p_{\text{Alice}}^* \rangle = \langle \Delta_0^+, e_{\mathcal{A}_0} p_{\text{Alice}}^* \rangle + \langle \Gamma_1^+, e_{\mathcal{A}_1} p_{\text{Alice}}^* \rangle.$$

As $\Gamma_1^+ \leq \Delta_1^+$ this quantity is at most

$$\langle \Delta_0^+, e_{\mathcal{A}_0} p_{\text{Alice}}^* \rangle + \langle \Delta_1^+, e_{\mathcal{A}_1} p_{\text{Alice}}^* \rangle.$$

Putting everything together, we have

$$\begin{aligned} \langle V, A_{\text{ns}} \otimes B \rangle &\leq \langle A, \Phi_V^*(B) \rangle + \langle \Delta_0^+, e_{\mathcal{A}_0} p_{\text{Alice}}^* \rangle + \langle \Delta_1^+, e_{\mathcal{A}_1} p_{\text{Alice}}^* \rangle \\ &= \langle A, \Phi_V^*(B) \rangle + \langle \text{mar}_{\mathcal{A}_1}(A) - A_0 \otimes e_{\mathcal{S}_1}^*, \Pi_0^A \rangle + \langle \text{mar}_{\mathcal{A}_0}(A) - A_1 \otimes e_{\mathcal{S}_0}^*, \Pi_1^A \rangle \\ &= \langle f_V(A, A_0, A_1), (B, \Pi_0^A, \Pi_1^A) \rangle. \end{aligned}$$

as desired. \square

Corollary 4.1 (Equivalence of min-max problems). *The following hold for any verifier V and any $\delta \geq 0$:*

1. $\mu(V) = \lambda(V)$.
2. If $(B^\mu, \Pi_0^\mu, \Pi_1^\mu)$ is δ -optimal for $\mu(V)$ then B^μ is δ -optimal for $\lambda(V)$.
3. If $(A^\mu, A_0^\mu, A_1^\mu)$ is δ -optimal for $\mu(V)$ then there exists A_{ns} such that A_{ns} is δ -optimal for $\lambda(V)$ and A_{ns} can be computed efficiently in parallel given $(A^\mu, A_0^\mu, A_1^\mu)$.

Proof. We begin with item 1. It is easy to prove $\lambda(V) \geq \mu(V)$: let A^λ be optimal for $\lambda(V)$, let A_0, A_1 witness the fact that A^λ is no-signaling, and let $(B^\mu, \Pi_0^\mu, \Pi_1^\mu)$ be optimal for $\mu(V)$. Then

$$\lambda(V) \geq \langle V, A^\lambda \otimes B^\mu \rangle = \langle f_V(A^\lambda, A_0, A_1), (B^\mu, \Pi_0^\mu, \Pi_1^\mu) \rangle \geq \mu(V).$$

For the reverse inequality, let $(A^\mu, A_0^\mu, A_1^\mu)$ be optimal for $\mu(V)$, let $\Pi_0^{A^\mu}, \Pi_1^{A^\mu}$ be optimal penalties for $(A^\mu, A_0^\mu, A_1^\mu)$, and let B^λ be optimal for $\lambda(V)$. By Theorem 4 there exists a no-signaling matrix A_{ns} witnessed by A_0^μ, A_1^μ such that

$$\langle V, A_{\text{ns}} \otimes B^\lambda \rangle \leq \langle f_V(A^\mu, A_0^\mu, A_1^\mu), (B^\lambda, \Pi_0^{A^\mu}, \Pi_1^{A^\mu}) \rangle.$$

The desired inequality $\lambda(V) \leq \mu(V)$ follows from the fact that the left side is at least $\lambda(V)$ and the right side is at most $\mu(V)$. The proof of item 1 is complete.

Item 2 follows easily from item 1. Let A be a no-signaling matrix and let A_0, A_1 witness this fact. Then

$$\lambda(V) - \delta = \mu(V) - \delta \leq \langle f_V(A, A_0, A_1), (B^\mu, \Pi_0^\mu, \Pi_1^\mu) \rangle = \langle V, A \otimes B^\mu \rangle.$$

As A was chosen arbitrarily, it follows that B^μ is δ -optimal for $\lambda(V)$.

For item 3, let B be any no-signaling matrix and let $\Pi_0^{A^\mu}, \Pi_1^{A^\mu}$ be optimal penalties for the given δ -optimal solution $(A^\mu, A_0^\mu, A_1^\mu)$. By Theorem 4 there exists a no-signaling matrix A_{ns} witnessed by A_0^μ, A_1^μ such that

$$\langle V, A_{\text{ns}} \otimes B \rangle \leq \langle f_V(A^\mu, A_0^\mu, A_1^\mu), (B, \Pi_0^{A^\mu}, \Pi_1^{A^\mu}) \rangle \leq \mu(V) + \delta = \lambda(V) + \delta.$$

As B was chosen arbitrarily, it follows that A_{ns} is δ -optimal for $\lambda(V)$. \square

3.4 Lemmas used in the rounding theorem

The lemmas used in the proof of Theorem 4 are not difficult. It is quite likely that some form of these lemmas is part of computer science “folklore,” though our notation may be nonstandard.

Lemma 5 (Small marginals have small preimages). *Let $a \in \mathcal{A}_{01}$ and $\vec{\delta} \in \mathcal{A}_0$ be nonnegative vectors with $\vec{\delta} \leq \text{mar}_{\mathcal{A}_1}(a)$. There exists a nonnegative vector $d \in \mathcal{A}_{01}$ with $d \leq a$ and $\text{mar}_{\mathcal{A}_1}(d) = \vec{\delta}$. Moreover, d can be computed efficiently in parallel given $a, \vec{\delta}$.*

Proof. Let $a_{(k_0, k_1)}$ and $\vec{\delta}_{k_0}$ denote the nonnegative entries of a and $\vec{\delta}$, respectively. Let s_{k_0} denote the k_0 th entry of $\text{mar}_{\mathcal{A}_1}(a)$ so that

$$s_{k_0} = \sum_{k_1=1}^{\dim(\mathcal{A}_1)} a_{(k_0, k_1)}.$$

The desired vector d has entries $d_{(k_0, k_1)}$ given by

$$d_{(k_0, k_1)} = \begin{cases} \vec{\delta}_{k_0} \frac{a_{(k_0, k_1)}}{s_{k_0}} & \text{when } s_{k_0} \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

(Intuitively, the weight $\vec{\delta}_{k_0}$ required of $\sum_{k_1} d_{(k_0, k_1)}$ is “spread out” over each $d_{(k_0, k_1)}$ proportionately according to $a_{(k_0, k_1)}$.) It is clear that this construction can be implemented efficiently in parallel.

Let us verify that $d \leq a$. Observe that for the case $s_{k_0} \neq 0$ the ratio $\vec{\delta}_{k_0}/s_{k_0}$ is at most one because $\vec{\delta} \leq \text{mar}_{\mathcal{A}_1}(a)$. Then

$$d_{(k_0, k_1)} = a_{(k_0, k_1)} \frac{\vec{\delta}_{k_0}}{s_{k_0}} \leq a_{(k_0, k_1)}$$

as desired. Of course, if $s_{k_0} = 0$ then $d_{(k_0, k_1)} = 0$ by definition and hence $d_{(k_0, k_1)} \leq a_{(k_0, k_1)}$ because $a \geq 0$.

Let us verify that $\text{mar}_{\mathcal{A}_1}(d) = \vec{\delta}$. For the case $s_{k_0} \neq 0$ the k_0 th entry of $\text{mar}_{\mathcal{A}_1}(d)$ is given by

$$\sum_{k_1=1}^{\dim(\mathcal{A}_1)} d_{(k_0, k_1)} = \frac{\vec{\delta}_{k_0}}{s_{k_0}} \sum_{k_1=1}^{\dim(\mathcal{A}_1)} a_{(k_0, k_1)} = \vec{\delta}_{k_0}$$

as desired. As above, if $s_{k_0} = 0$ then by definition $d_{(k_0, k_1)} = 0$ for each k_1 and hence $\sum_{k_1} d_{(k_0, k_1)} = 0$. As $0 \leq \vec{\delta}_{k_0} \leq s_{k_0}$ it must be that $\vec{\delta}_{k_0} = 0$, too. \square

Corollary 5.1. *Let $A : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$ and $\Delta : \mathcal{S}_{01} \rightarrow \mathcal{A}_0$ be nonnegative matrices with $\Delta \leq \text{mar}_{\mathcal{A}_1}(A)$. There exists a nonnegative matrix $D : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$ with $D \leq A$ and $\text{mar}_{\mathcal{A}_1}(D) = \Delta$. Moreover, D can be computed efficiently in parallel given A, Δ .*

Proof. Apply Lemma 5 to each of the columns of A, Δ . \square

Lemma 6 (Disjoint marginals are always consistent). *For both $c \in \{0, 1\}$ let $t_c \in \mathcal{A}_c$ be nonnegative vectors whose entries sum to the same value. There exists a nonnegative vector $t \in \mathcal{A}_{01}$ with $\text{mar}_{\mathcal{A}_c}(t) = t_c$ for both $c \in \{0, 1\}$. Moreover, t can be computed efficiently in parallel given t_0, t_1 .*

Proof. Let p_{k_0} and q_{k_1} be the nonnegative entries of t_0 and t_1 , respectively. Let s denote the sum of the entries of t_0, t_1 so that

$$s = \sum_{k_0=1}^{\dim(\mathcal{A}_0)} p_{k_0} = \sum_{k_1=1}^{\dim(\mathcal{A}_1)} q_{k_1}.$$

If $s = 0$ then it is clear that the desired vector t is the zero vector. For the remainder of the proof assume that $s \neq 0$. The desired vector t has entries $t_{(k_0, k_1)}$ given by

$$t_{(k_0, k_1)} = \frac{p_{k_0} q_{k_1}}{s}$$

It is clear that this construction can be implemented efficiently in parallel.

Let us verify that $\text{mar}_{\mathcal{A}_c}(t) = t_c$ for both $c \in \{0, 1\}$. For the case $c = 0$ the k_0 th entry of $\text{mar}_{\mathcal{A}_1}(t)$ is given by

$$\sum_{k_1=1}^{\dim(\mathcal{A}_1)} \frac{p_{k_0} q_{k_1}}{s} = \frac{p_{i,j} s}{s} = p_{k_0}$$

as desired. The case $c = 1$ is handled similarly. \square

Corollary 6.1. *For both $c \in \{0, 1\}$ let $T_c : \mathcal{S}_{01} \rightarrow \mathcal{A}_c$ be nonnegative matrices with $\text{mar}_{\mathcal{A}_0}(T_0) = \text{mar}_{\mathcal{A}_1}(T_1)$. There exists a nonnegative matrix $T : \mathcal{S}_{01} \rightarrow \mathcal{A}_{01}$ with $\text{mar}_{\mathcal{A}_c}(T) = T_c$ for both $c \in \{0, 1\}$. Moreover, T can be computed efficiently in parallel given T_0, T_1 .*

Proof. Apply Lemma 6 to each of the columns of T_0, T_1 . \square

4 A parallel multiplicative weights algorithm

In this section we complete the proof of our main result—that every decision problem that admits a two-turn interactive proof with competing teams of no-signaling provers is also in PSPACE. Most of the detail appears in Section 4.1 wherein we present an efficient parallel oracle-algorithm based on the MWUM that produces δ -optimal no-signaling strategies for the teams, given an oracle for “best responses” for Team Bob to a given candidate strategy for Alice. We describe an efficient parallel implementation of the required oracle in Section 4.2, from which the unconditional efficiency of our algorithm immediately follows. The ensuing inclusion of $\text{MRG}_{\text{ns}}(2, 2)$ inside PSPACE is discussed in Section 4.3.

4.1 The parallel algorithm

Precise statements of the problem solved by our algorithm and the oracle it requires are given below. All input numbers are written as rational numbers in binary. For matrix inputs, each entry is written explicitly.

Problem 1 (Weak no-signaling equilibrium).

Input: A verifier matrix $V : \mathcal{S}_{01} \mathcal{T}_{01} \rightarrow \mathcal{A}_{01} \mathcal{B}_{01}$ and an accuracy parameter $\delta > 0$.

Oracle: Weak no-signaling optimization. (See Problem 2 below.)

Output: δ -optimal no-signaling strategies \tilde{A}, \tilde{B} for the min-max problem $\lambda(V)$.

Problem 2 (Weak no-signaling optimization).

Input: A verifier-Alice matrix $S : \mathcal{T}_{01} \rightarrow \mathcal{B}_{01}$ and an accuracy parameter $\delta > 0$.

Output: A δ -optimal no-signaling strategy \tilde{B} for Team Bob. (That is, a no-signaling matrix \tilde{B} such that $\langle S, \tilde{B} \rangle \geq \langle S, B \rangle - \delta$ for all no-signaling matrices B .)

Given Corollary 4.1, it suffices to find δ -optimal solutions $(\tilde{A}, \tilde{A}_0, \tilde{A}_1)$ and $(\tilde{B}, \tilde{\Pi}_0, \tilde{\Pi}_1)$ for $\mu(V)$ and then convert these solutions into δ -optimal strategies for $\lambda(V)$. This method is codified in the algorithm of Figure 2.

This algorithm is a straightforward modification of the standard multiplicative weights update method for equilibrium problems. The precise formulation of the MWUM used in this paper is stated as Theorem 7. Our statement of this theorem is somewhat nonstandard: the result is usually presented in the form of an algorithm, whereas our presentation is purely mathematical. However, a cursory examination of the literature—say, Kale’s thesis [Kal07, Chapter 2]—reveals that our mathematical formulation is equivalent to the more conventional algorithmic form.

Theorem 7 (Multiplicative weights update method—see Ref. [Kal07, Theorem 2]). *Fix an $\varepsilon \in (0, 1/2)$. Let m^1, \dots, m^T be arbitrary D -dimensional “loss” vectors whose entries m_i^t lay in the interval $[-\alpha, \alpha]$. Let w^1, \dots, w^T be D -dimensional nonnegative “weight” vectors whose entries w_i^t are given recursively via*

$$\begin{aligned} w_i^1 &= 1 \\ w_i^{t+1} &= w_i^t (1 - \varepsilon m_i^t). \end{aligned}$$

Let p^1, \dots, p^T be probability vectors obtained by normalizing each w^1, \dots, w^T . For all probability vectors p it holds that

$$\frac{1}{T} \sum_{t=1}^T \langle p^t, m^t \rangle \leq \left\langle p, \frac{1}{T} \sum_{t=1}^T m^t \right\rangle + \alpha \left(\varepsilon + \frac{\ln D}{\varepsilon T} \right).$$

Note that Theorem 7 holds for *all* choices of loss vectors m^1, \dots, m^T , including the case in which each m^t is chosen adversarially based upon w^t . This adaptive selection of loss vectors is typical in implementations of the MWUM.

Proposition 8. *The oracle-algorithm presented in Figure 2 solves the weak no-signaling equilibrium problem (Problem 1). Assuming unit cost for the oracle, this algorithm can be implemented in parallel with run time bounded by a polynomial in $1/\delta$ and $\log(\dim(\mathcal{S}_{01}\mathcal{T}_{01}\mathcal{A}_{01}\mathcal{B}_{01}))$.*

Proof. For each pair $i = (i_0, i_1)$ of questions let π_i denote the probability with which the verifier asks questions i to Team Alice. Let m^t denote the i th column of M^t for each $t = 1, \dots, T$. We argue that the entries of m^t lay in the interval $[0, 3\pi_i]$. To this end, observe that the loss matrix M^t is defined in Figure 2 via the adjoint mapping f_V^* as

$$M^t = \Phi_V^*(B^t) + e_{\mathcal{A}_1} \otimes \Pi_0^t + e_{\mathcal{A}_0} \otimes \Pi_1^t \leq 3e_{\mathcal{A}_{01}} p_{\text{Alice}}^*$$

where the inequality follows immediately from the bound $\Phi_V^*(B) \leq e_{\mathcal{A}_{01}} p_{\text{Alice}}^*$ of Proposition 3 and the restriction $\Pi_c \leq e_{\mathcal{A}_c} p_{\text{Alice}}^*$ on penalty matrices. The desired bound on the entries of m^t follows from the observation that the i th column of $3e_{\mathcal{A}_{01}} p_{\text{Alice}}^*$ is the vector whose entries are all equal to $3\pi_i$.

1. Let $\varepsilon = \delta/10$ and let $T = \left\lceil \frac{\ln(\dim(\mathcal{A}_{01}))}{\varepsilon^2} \right\rceil$.

Let (W^1, W_0^1, W_1^1) denote the triple of all-ones matrices and let (A^1, A_0^1, A_1^1) denote the uniformly random strategy for Alice obtained by normalizing the columns of (W^1, W_0^1, W_1^1) .

2. Repeat for each $t = 1, \dots, T$:

- (a) Compute optimal penalties Π_0^t, Π_1^t for (A^t, A_0^t, A_1^t) as described in Section 3.2. Use the oracle for Problem 2 to obtain a $\delta/2$ -best response B^t to the verifier-Alice matrix $\Phi_V(A^t)$.
- (b) Compute the loss matrices $(M^t, M_0^t, M_1^t) = f_V^*(B^t, \Pi_0^t, \Pi_1^t)$. Exit the loop now if $t = T$.
- (c) Update the weight matrices according to the standard multiplicative weights update rule:

$$(W^{t+1}, W_0^{t+1}, W_1^{t+1}) = (W^t, W_0^t, W_1^t) \boxtimes \left(\underbrace{(W^1, W_0^1, W_1^1)}_{\text{all-ones matrices}} - \varepsilon (M^t, M_0^t, M_1^t) \right)$$

where \boxtimes denotes the (entrywise) matrix Schur product. (See Theorem 7.)

- (d) Compute the updated triple $(A^{t+1}, A_0^{t+1}, A_1^{t+1})$ of stochastic matrices for Team Alice by normalizing the columns of $(W^{t+1}, W_0^{t+1}, W_1^{t+1})$.

3. Compute

$$(\tilde{A}, \tilde{A}_0, \tilde{A}_1) = \frac{1}{T} \sum_{t=1}^T (A^t, A_0^t, A_1^t) \quad \text{and} \quad (\tilde{B}, \tilde{\Pi}_0, \tilde{\Pi}_1) = \frac{1}{T} \sum_{t=1}^T (B^t, \Pi_0^t, \Pi_1^t)$$

both of which are δ -optimal for $\mu(V)$. Compute the no-signaling matrix \tilde{A}_{ns} from $(\tilde{A}, \tilde{A}_0, \tilde{A}_1)$ as described in Corollary 4.1.

4. Return $(\tilde{A}_{\text{ns}}, \tilde{B})$ as the δ -optimal strategies of Team Alice and Team Bob for $\lambda(V)$.

Figure 2: Algorithm that finds δ -optimal solutions to the equilibrium problem $\lambda(V)$ for two-turn interactive proofs with competing teams of no-signaling provers (Problem 1).

Let a^t denote the i th column of A^t for $t = 1, \dots, T$. It is clear that the construction of the probability vectors a^t in terms of the loss vectors m^t presented in Figure 2 obeys the condition of Theorem 7. It therefore follows that for any probability vector $a \in \mathcal{A}_{01}$ we have

$$\frac{1}{T} \sum_{t=1}^T \langle a^t, m^t \rangle \leq \left\langle a, \frac{1}{T} \sum_{t=1}^T m^t \right\rangle + 3\pi_i \left(\varepsilon + \frac{\ln(\dim(\mathcal{A}_{01}))}{\varepsilon T} \right).$$

Summing these inequalities over all columns i we find that for any stochastic matrix A it holds that

$$\frac{1}{T} \sum_{t=1}^T \langle A^t, M^t \rangle \leq \left\langle A, \frac{1}{T} \sum_{t=1}^T M^t \right\rangle + 3 \left(\varepsilon + \frac{\ln(\dim(\mathcal{A}_{01}))}{\varepsilon T} \right).$$

A similar bound on the stochastic matrices A_0^t, A_1^t in terms of the loss matrices M_0^t, M_1^t can be derived in much the same way. For completeness, let us make this argument explicit. For both $c \in \{0, 1\}$ and for each question i_c let π_{i_c} denote the probability with which the referee asks question i_c to Alice _{c} . Let m_c^t denote the i_c th column of M_c^t for each $t = 1, \dots, T$. We argue that the entries of m_c^t lay in the interval $[-\pi_{i_c}, 0]$. Recall the loss matrix M_c^t is defined in Figure 2 via the adjoint mapping f_V^* as

$$M_c^t = -\Pi_c^t (I_{\mathcal{S}_c} \otimes e_{\mathcal{S}_{\bar{c}}}) \geq -e_{\mathcal{A}_c} \text{mar}_{\mathcal{S}_{\bar{c}}}(p_{\text{Alice}})^*$$

where the inequality follows immediately from the restriction $\Pi_c \leq e_{\mathcal{A}_c} p_{\text{Alice}}^*$ on penalty matrices. The desired bound on the entries of m_c^t follows from the observation that the i_c th column of $e_{\mathcal{A}_c} \text{mar}_{\mathcal{S}_{\bar{c}}}(p_{\text{Alice}})^*$ is the vector whose entries are all equal to π_{i_c} .

As above, let a_c^t denote the i_c th column of A_c^t for $t = 1, \dots, T$. It is clear that the construction of the probability vectors a_c^t in terms of the loss vectors m_c^t presented in Figure 2 obeys the condition of Theorem 7. It therefore follows that for any probability vector $a_c \in \mathcal{A}_c$ we have

$$\frac{1}{T} \sum_{t=1}^T \langle a_c^t, m^t \rangle \leq \left\langle a_c, \frac{1}{T} \sum_{t=1}^T m_c^t \right\rangle + \pi_{i_c} \left(\varepsilon + \frac{\ln(\dim(\mathcal{A}_c))}{\varepsilon T} \right).$$

Summing these inequalities over all columns i_c we find that for any stochastic matrix A_c it holds that

$$\frac{1}{T} \sum_{t=1}^T \langle A_c^t, M_c^t \rangle \leq \left\langle A_c, \frac{1}{T} \sum_{t=1}^T M_c^t \right\rangle + \varepsilon + \frac{\ln(\dim(\mathcal{A}_c))}{\varepsilon T}.$$

At this point we have derived three inequalities for three arbitrary stochastic matrices A, A_0, A_1 . Summing these inequalities and substituting $(M^t, M_0^t, M_1^t) = f_V^*(B^t, \Pi_0^t, \Pi_1^t)$ and the choices of ε, T listed in Figure 2 we find that for any triple (A, A_0, A_1) of stochastic matrices it holds that

$$\frac{1}{T} \sum_{t=1}^T \langle f_V(A^t, A_0^t, A_1^t), (B^t, \Pi_0^t, \Pi_1^t) \rangle \leq \left\langle f_V(A, A_0, A_1), \frac{1}{T} \sum_{t=1}^T (B^t, \Pi_0^t, \Pi_1^t) \right\rangle + \delta/2. \quad (1)$$

The remainder of this proof is a straightforward adaptation of Kale's analysis for the much simpler class of two-player zero-sum games in normal form [Kal07, Section 2.3.1]. We argue that the triples $(\hat{A}, \hat{A}_0, \hat{A}_1)$

and $(\tilde{B}, \tilde{\Pi}_0, \tilde{\Pi}_1)$ appearing Figure 2 are δ -optimal for $\mu(V)$. Let us begin with the triple $(\tilde{A}, \tilde{A}_0, \tilde{A}_1)$. Choose any (B, Π_0, Π_1) and let (A^*, A_0^*, A_1^*) be optimal for $\mu(V)$. We have

$$\begin{aligned} \left\langle \frac{1}{T} \sum_{t=1}^T f_V(A^t, A_0^t, A_1^t), (B, \Pi_0, \Pi_1) \right\rangle &\leq \frac{1}{T} \sum_{t=1}^T \langle f_V(A^t, A_0^t, A_1^t), (B^t, \Pi_0^t, \Pi_1^t) \rangle + \delta/2 \\ &\leq \left\langle f_V(A^*, A_0^*, A_1^*), \frac{1}{T} \sum_{t=1}^T (B^t, \Pi_0^t, \Pi_1^t) \right\rangle + \delta \leq \mu(V) + \delta \end{aligned}$$

as desired. (The first inequality is because each (B^t, Π_0^t, Π_1^t) is a $\delta/2$ -best response to (A^t, A_0^t, A_1^t) ; the second is Eq. (1).)

To see that $(\tilde{B}, \tilde{\Pi}_0, \tilde{\Pi}_1)$ is δ -optimal for $\mu(V)$, let (A, A_0, A_1) be any triple of stochastic matrices. We have

$$\left\langle f_V(A, A_0, A_1), \frac{1}{T} \sum_{t=1}^T (B^t, \Pi_0^t, \Pi_1^t) \right\rangle \geq \frac{1}{T} \sum_{t=1}^T \langle f_V(A^t, A_0^t, A_1^t), (B^t, \Pi_0^t, \Pi_1^t) \rangle - \delta/2 \geq \mu(V) - \delta$$

as desired. (The first inequality is Eq. (1); the second is because each (B^t, Π_0^t, Π_1^t) is a $\delta/2$ -best response to (A^t, A_0^t, A_1^t) .) Finally, it follows from Corollary 4.1 that \tilde{A}_{ns} and \tilde{B} are δ -optimal strategies for $\lambda(V)$.

That the algorithm admits an efficient parallel implementation is straightforward. In each iteration computations of optimal penalties, the loss matrices (via f_V^*), the multiplicative weights update rule, and normalization are all simple operations involving only addition and multiplication of individual rational entries of matrices that can easily be implemented in parallel. Efficiency follows from the fact that the total number of iterations is bounded by a polynomial in $1/\delta$ and the logarithm of $\dim(\mathcal{S}_{01}\mathcal{T}_{01}\mathcal{A}_{01}\mathcal{B}_{01})$, the size of the verifier matrix. \square

4.2 Implementations of the best-response oracle for Team Bob

In order for the algorithm of Figure 2 to be unconditionally efficient, we require a parallel implementation of the oracle for weak no-signaling optimization (Problem 2). Fortunately, all the work is already done: Problem 2 is the optimization problem that arises naturally from two-turn, two-prover interactive proofs with no-signaling provers. Thus, the parallel algorithm of Ito [Ito10] can be re-used to implement the oracle in our algorithm without complication.

In Ito's terminology, the verifier-Alice matrix $\Phi_V(A)$ specifies a *game* and the two no-signaling provers comprising Team Bob are the *players*. Ito does not claim that an explicit strategy for the players can be found efficiently in parallel. Rather, he claims only that the task of distinguishing high success probability from low success probability admits a parallel algorithm, as this simpler task is sufficient to put $\text{MIP}_{\text{ns}}(2, 2)$ inside PSPACE. However, a cursory glance at the details of Ito's proof reveals a parallel construction of near-optimal no-signaling strategies for the players as required by Problem 2.

Alternatively, the oracle for weak no-signaling optimization (Problem 2) can be implemented by re-using the algorithm for weak no-signaling equilibrium (Problem 1) listed in Figure 2 of the present paper. Indeed, Problem 2 is a special case of Problem 1 in which one team has a trivial strategy space. In this special case the required "oracle" demands only weak no-signaling optimization over a trivial strategy space, which of course admits a trivial parallel implementation. In other words, the algorithm of Figure 2 can be used in a two-level recursive fashion to give an unconditionally efficient parallel algorithm for Problem 1.

4.3 Containment in PSPACE

The desired containment of $\text{MRG}_{\text{ns}}(2, 2)$ inside PSPACE now follows in the usual way:

Theorem 1. *Every decision problem that admits a two-turn interactive proof with competing teams of two no-signaling provers per team is also in PSPACE. Thus, we obtain the identity $\text{MRG}_{\text{ns}}(2, 2) = \text{PSPACE}$.*

Proof. Let L be a decision problem in $\text{MRG}_{\text{ns}}(2, 2)$ with completeness c and soundness s and let x be any input string. Each entry of the exponential-size verifier matrix $V : \mathcal{S}_{01}\mathcal{T}_{01} \rightarrow \mathcal{A}_{01}\mathcal{B}_{01}$ induced by the verifier on input x can be computed in space polynomial in $|x|$ by simulating every choice of randomness for the verifier. In order to decide whether x is a yes-instance or no-instance of L it suffices to find δ -optimal strategies for the teams for $\delta = (c - s)/3$, which permits us to distinguish $\lambda(V) \geq c$ from $\lambda(V) \leq s$. It follows from Proposition 8 and the discussion in Section 4.2 that the algorithm of Figure 2 can be used to find δ -optimal strategies for the teams and can be implemented in parallel with run time bounded by a polynomial in $1/\delta$ and the logarithm of the dimensions of V . As the dimensions of V scale exponentially with $|x|$ and δ scales as an inverse polynomial in $|x|$ the total run time of this parallel algorithm scales polynomially with $|x|$ and can therefore be simulated in polynomial space in the usual way [Bor77]. \square

Acknowledgements

The author is grateful to Tsuyoshi Ito, Sarvagya Upadhyay, John Watrous, and Xiaodi Wu for helpful discussions. This research is supported by the Government of Canada through Industry Canada, the Province of Ontario through the Ministry of Research and Innovation, NSERC, DTO-ARO, CIFAR, and Quantum-Works.

References

- [AHK05] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta algorithm and applications. Submitted, 2005.
- [Bab85] László Babai. Trading group theory for randomness. In *Proceedings of the 17th ACM Symposium on Theory of Computing (STOC 1985)*, pages 421–429, 1985.
- [BBL⁺06] Gilles Brassard, Harry Buhrman, Noah Linden, André Méthot, Alain Tapp, and Falk Unger. A limit on nonlocality in any world in which communication complexity is not trivial. *Physical Review Letters*, 96(25):250401, 2006. arXiv:quant-ph/0508042v1.
- [Bel64] John Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1(3):195–200, 1964.
- [BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1(1):3–40, 1991.
- [BM88] László Babai and Shlomo Moran. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36(2):254–276, 1988.
- [BOGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: how to remove intractability assumptions. In *Proceedings of the 20th ACM Symposium on Theory of Computing (STOC 1988)*, pages 113–131, 1988.

- [Bor77] Allan Borodin. On relating time and space to size and depth. *SIAM Journal on Computing*, 6(4):733–744, 1977.
- [CCL94] Jin-Yi Cai, Anne Condon, and Richard Lipton. PSPACE is provable by two provers in one round. *Journal of Computer and System Sciences*, 48(1):183–193, 1994.
- [CGJ09] Richard Cleve, Dmitry Gavinsky, and Rahul Jain. Entanglement-resistant two-prover interactive proof systems and non-adaptive PIRs. *Quantum Information and Computation*, 9:648–656, 2009. arXiv:0707.1729v1 [quant-ph].
- [CHTW04] Richard Cleve, Peter Høyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th Conference on Computational Complexity*, pages 236–249, 2004. arXiv:quant-ph/0404076v1.
- [Fan53] K. Fan. Minimax theorems. *Proceedings of the National Academy of Sciences*, 39:42–47, 1953.
- [FK97] Uriel Feige and Joe Kilian. Making games short. In *Proceedings of the 29th ACM Symposium on Theory of Computing (STOC 1997)*, pages 506–516, 1997.
- [FKS95] Joan Feigenbaum, Daphne Koller, and Peter Shor. A game-theoretic classification of interactive complexity classes. In *Proceedings of the 10th Conference on Structure in Complexity Theory*, pages 227–237, 1995.
- [FL92] Uriel Feige and László Lovász. Two-prover one-round proof systems: their power and their problems. In *Proceedings of the 24th ACM Symposium on Theory of Computing (STOC 1992)*, pages 733–744, 1992.
- [FRS94] Lance Fortnow, John Rempel, and Michael Sipser. On the power of multi-prover interactive protocols. *Theoretical Computer Science*, 134(2):545–557, 1994.
- [FS92] Uriel Feige and Adi Shamir. Multi-oracle interactive protocols with constant space verifiers. *Journal of Computer and System Sciences*, 44:259–271, 1992.
- [FST90] Uriel Feige, Adi Shamir, and Moshe Tennenholtz. The noisy oracle problem. In *Advances in Cryptology – Proceedings of Crypto’88*, volume 403 of *Lecture Notes in Computer Science*, pages 284–296. Springer, 1990.
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [GS89] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In Silvio Micali, editor, *Randomness and Computation*, volume 5 of *Advances in Computing Research*, pages 73–90. JAI Press, 1989.
- [GW11] Gus Gutoski and Xiaodi Wu. Parallel approximation of min-max problems with applications to classical and quantum zero-sum games. arXiv:1011.2787v2 [quant-ph], 2011.
- [Hol09] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5:141–172, 2009. arXiv:cs/0607139v3 [cs.CC].

- [IKM09] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *Proceedings of the 24th IEEE Conference on Computational Complexity (CCC 2009)*, pages 217–228, 2009. arXiv:0810.0693v1 [quant-ph].
- [Ito10] Tsuyoshi Ito. Polynomial-space approximation of no-signaling provers. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP 2010)*, volume 6198 of *Lecture Notes in Computer Science*, pages 140–151. Springer, 2010. arXiv:0908.2363v2 [cs.CC].
- [JJUW10] Rahul Jain, Zhengfeng Ji, Sarvagya Upadhyay, and John Watrous. QIP=PSPACE. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC 2010)*, pages 573–582, 2010. arXiv:0907.4737v2 [quant-ph].
- [JUW09] Rahul Jain, Sarvagya Upadhyay, and John Watrous. Two-message quantum interactive proofs are in PSPACE. In *Proceedings of the 50th IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, pages 534–543, 2009. arXiv:0905.1300v1 [quant-ph].
- [JW09] Rahul Jain and John Watrous. Parallel approximation of non-interactive zero-sum quantum games. In *Proceedings of the 24th IEEE Conference on Computational Complexity (CCC 2009)*, pages 243–253, 2009. arXiv:0808.2775v1 [quant-ph].
- [Kal07] Satyen Kale. *Efficient algorithms using the multiplicative weights update method*. PhD thesis, Princeton University, 2007.
- [KM92] Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4:528–552, 1992.
- [KMvS94] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Fast algorithms for finding randomized strategies in game trees. In *Proceedings of the 26th ACM Symposium on Theory of Computing (STOC 1994)*, pages 750–759, 1994.
- [KR10] Julia Kempe and Oded Regev. No strong parallel repetition with entangled and non-signaling provers. In *Proceedings of the 25th IEEE Conference on Computational Complexity (CCC 2010)*, pages 7–15, 2010. arXiv:0911.0201v1 [quant-ph].
- [LFKN92] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *Journal of the ACM*, 39(4):859–868, 1992.
- [Sha92] Adi Shamir. IP = PSPACE. *Journal of the ACM*, 39(4):869–877, 1992.
- [vD05] Wim van Dam. Implausible consequences of superstrong nonlocality. arXiv:quant-ph/0501159v1, 2005.
- [Vil38] Jean Ville. Sur la théorie générale des jeux où intervient l’habileté des joueurs. *Traité du calcul des probabilités et des applications*, IV(2):105–113, 1938. In French.
- [WK06] Manfred Warmuth and Dima Kuzmin. Online variance minimization. In *Proceedings of the 19th Annual Conference on Learning Theory*, volume 4505 of *Lecture Notes in Computer Science*, pages 514–528. Springer, 2006.

- [Wu10] Xiaodi Wu. Equilibrium value method for the proof of $\text{QIP}=\text{PSPACE}$. arXiv:1004.0264v2 [quant-ph], 2010.
- [You01] Neal Young. Sequential and parallel algorithms for mixed packing and covering. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS 2001)*, pages 538–546, 2001.